

Frequency Stability Review

C. A. Greenhall

Communications Systems Research Section

This tutorial article treats certain aspects of the description and measurement of oscillator stability. Topics covered are time and frequency deviations, Allan variance, the zero-crossing counter measurement technique, frequency drift removal, and the "three-cornered hat."

I. Deviations in Phase, Time, and Frequency

The purpose of this article is to define the Allan variance, relate it to various kinds of disturbances in oscillators and measurement systems, and to describe, from the point of view of an analysis and software person, the frequency stability measurement method used at the JPL Frequency Standards Laboratory. This is an expanded version of a talk given in Oct 1983 to DSN operations. It outlines the JPL frequency stability measurement methodology.

Let oscillators 1 and 2, both running at nominal frequency ν_0 , have the outputs

$$\sin(2\pi\nu_0 t + \phi_1(t)), \quad \sin(2\pi\nu_0 t + \phi_2(t))$$

where $\phi_1(t)$ and $\phi_2(t)$ are the phase deviations relative to frequency ν_0 . Suppose that both oscillators are used for driving clocks. At true time t , the clock times t_1, t_2 are defined by

$$2\pi\nu_0 t_1 = 2\pi\nu_0 t + \phi_1(t) \quad (1)$$

$$2\pi\nu_0 t_2 = 2\pi\nu_0 t + \phi_2(t) \quad (2)$$

In other words, t_i is the time at which the output of a perfect oscillator, $\sin(2\pi\nu_0 t)$, would have the same total phase as oscillator i at time t .

The relative phase and time deviations of the pair of oscillators are defined by

$$\phi(t) = \phi_1(t) - \phi_2(t)$$

$$x(t) = t_1 - t_2$$

Subtracting (2) from (1) gives

$$x(t) = \frac{\phi(t)}{2\pi\nu_0}$$

The fractional frequency deviation is defined by

$$y(t) = \frac{1}{2\pi\nu_0} \frac{d\phi}{dt}$$

Then

$$y(t) = \frac{dx}{dt}$$

$$x(t) = \int y(t) dt$$

The average fractional frequency deviation over the time interval $(t-\tau, t)$ is defined by

$$\bar{y}(t, \tau) = \frac{1}{\tau} \int_{t-\tau}^t y(s) ds$$

$$= \frac{x(t) - x(t-\tau)}{\tau}$$

$$= \frac{\Delta_{\tau} x(t)}{\tau}$$

where Δ_{τ} is the backward first difference operator. We call τ the averaging time. When someone says that two oscillators differ in frequency by 5×10^{-14} , say, he is most likely quoting a value of $\bar{y}(t, \tau)$. To understand the quote, the value of τ must be known. Is it one second or one day?

Figure 1 shows plots of $x(t)$ and $\bar{y}(t, \tau)$ from a test of two hydrogen masers, where $\tau = 4094$ s and both x and \bar{y} are sampled once every τ . The lower curve \bar{y} is approximately the derivative of the upper curve x . Notice that $x = 0$ at both endpoints. This is an artifact of the presentation; the mean \bar{y} has been subtracted off.

II. Allan Variance

The Allan variance was invented to solve the problem of characterizing the RMS frequency deviations of a pair of oscillators. Given time-deviation samples $x(0), x(\tau), \dots, x(m\tau)$ over the time $T = m\tau$, let

$$\bar{y}_j = \bar{y}(j\tau, \tau)$$

the average frequency between $(j-1)\tau$ and $j\tau$. One could compute the averages

$$\langle x \rangle = \frac{1}{m+1} \sum_{j=0}^m x(j\tau)$$

$$\langle \bar{y} \rangle = \frac{1}{m} \sum_{j=1}^m \bar{y}_j$$

and the RMS deviations

$$x_{\text{RMS}} = \left[\frac{1}{m+1} \sum_{j=0}^m (x(j\tau) - \langle x \rangle)^2 \right]^{1/2}$$

$$\bar{y}_{\text{RMS}} = \left[\frac{1}{m} \sum_{j=1}^m (\bar{y}_j - \langle \bar{y} \rangle)^2 \right]^{1/2}$$

Often, however, these quantities tend to grow with T ; moreover, they may fail to be reproducible from one T -interval to the next. (For example, if \bar{y} is a random walk, then the mean and standard deviation of \bar{y}_{RMS}^2 are both proportional to T .)

For many types of oscillator phase deviations, the Allan variance, denoted by $\sigma_y^2(\tau)$, avoids these problems while still yielding a meaningful measure of frequency stability. To define theoretical Allan variance from the preceding setup, we must assume that $x(j\tau)$ is available for all non-negative integers j . Then, by definition,

$$\sigma_y^2(\tau) = \lim_{m \rightarrow \infty} \frac{1}{2(m-1)} \sum_{j=2}^m (\bar{y}_j - \bar{y}_{j-1})^2 \quad (3)$$

if the limit exists. In other words, $\sigma_y(\tau)$ is the RMS average of the quantities $(\bar{y}_j - \bar{y}_{j-1})/\sqrt{2}$; thus $\sigma_y(\tau)$ measures the RMS change in $\bar{y}(t, \tau)$ when t changes by τ . Note that

$$\bar{y}_j - \bar{y}_{j-1} = \frac{1}{\tau} [\Delta_{\tau} x(j\tau) - \Delta_{\tau} x((j-1)\tau)]$$

$$= \frac{1}{\tau} \Delta_{\tau}^2 x(j\tau)$$

$$= \frac{1}{\tau} [x(j\tau) - 2x((j-1)\tau) + x((j-2)\tau)]$$

where Δ_{τ}^2 is the backward second difference operator. The quantity $\sigma_y^2(\tau)$ is called the Allan variance for the averaging time τ . It seems that $\sigma_y(\tau)$ (without the square) should be called Allan deviation. In this writeup, it will be done so for the sake of precision. In casual discussions, however, $\sigma_y(\tau)$ is called the Allan variance without confusing anyone.

The factor 2 (or $\sqrt{2}$) is there for historical reasons. On the other hand, if the two oscillators are judged to be of like quality, then the Allan deviation of the pair is sometimes divided by $\sqrt{2}$ to give the Allan deviation of the individual oscillators. This is a different $\sqrt{2}$. Moreover, this procedure assumes that the second τ -differences of phase of the two oscillators are

orthogonal random processes. If linear frequency drift is present, or if the two oscillators are subject to the same environmental fluctuations, then this assumption is unrealistic, at least for large τ .

The usual estimator of $\sigma_y^2(\tau)$, given $x(t)$ for $0 \leq t \leq T = m\tau$, is

$$\begin{aligned} S_y^2(\tau, m) &= \frac{1}{2(m-1)} \sum_{j=2}^m (\bar{y}_j - \bar{y}_{j-1})^2 \\ &= \frac{1}{2\tau^2(m-1)} \sum_{j=2}^m [\Delta_\tau^2 x(j\tau)]^2 \end{aligned} \quad (4)$$

Table 1 gives a numerical example, a fragment from an actual test of two hydrogen masers. Elapsed time t is given in units of $\tau_0 = 256$ s, and $x(t)$ is given in units of 10^{-14} s. From these data, the following estimates can be calculated:

$$\begin{aligned} S_y(\tau_0, 8) &= \frac{10^{-14}}{\sqrt{2} \tau_0} \left[\frac{1}{7} (87^2 + \dots + 7^2) \right]^{1/2} \\ &= 2.92 \times 10^{-15} \\ S_y(2\tau_0, 4) &= \frac{10^{-14}}{\sqrt{2} 2\tau_0} \left[\frac{1}{3} (125^2 + 56^2 + 37^2) \right]^{1/2} \\ &= 1.13 \times 10^{-15} \\ S_y(3\tau_0, 2) &= \frac{10^{-14}}{\sqrt{2} 3\tau_0} 91 \\ &= 8.37 \times 10^{-16} \end{aligned}$$

There are other estimators of $\sigma_y(\tau)$. The RMS average of the second differences of x can include values of $\Delta_\tau^2 x(t)$ for t not a multiple of τ . In the last example, for $\tau = 3\tau_0$, the values $\Delta_\tau^2 x(t)$, $t = 6, 7, 8$, can be used for computing the estimate

$$\frac{10^{-14}}{\sqrt{2} 3\tau_0} \left[\frac{1}{3} (91^2 + 87^2 + 63^2) \right]^{1/2} = 7.48 \times 10^{-16}$$

Is this a better estimate than $S_y(3\tau_0, 2)$? If so, then in what sense is it better? This subject is controversial (Refs. 5, 6, and 8).¹

¹Also, D. Percival, letter to J. Barnes, Aug. 31, 1982.

Figure 2 shows $S_y(\tau, m)$ ("sigma") versus τ for the same test that yielded Fig. 1.

A. Naive Error Estimate

If, in the last sum of (4), the $\Delta_\tau^2 x(j\tau)$ were independent, zero-mean Gaussian random variables with the same variance, then $S_y^2(\tau, m)$ would be proportional to a χ^2 variable with $m-1$ degrees of freedom. The assumptions of independence and zero mean are almost never realistic. Not knowing the true situation, however, one usually estimates the standard deviation of $S_y^2(\tau, m)$ by the naive formula $S_y^2(\tau, m)\epsilon$, where $\epsilon = \sqrt{2/(m-1)}$. Then

$$\left[S_y(\tau, m) \sqrt{1-\epsilon}, S_y(\tau, m) \sqrt{1+\epsilon} \right] \quad (5)$$

is presented as the roughest sort of "one-sigma" error bar for $S_y(\tau, m)$. If drift is removed (Section IV), then $\epsilon = \sqrt{2/(m-2)}$ is used. Figures 2 and 6 show the intervals (5) and also give $m-1$, the "number of samples."

Don't use (5) unless $m \geq 4$; even then, don't take it seriously. Naturally, if you know the properties of the process $x(t)$ in advance, you can compute more accurate variance estimators and confidence intervals (Refs. 2, 3, 5, 6, 7, and 8).¹ But if you did know all this, you wouldn't be testing the oscillators in the first place. In any case, make sure that the user of the results knows the number of samples.

B. Deterministic Examples

The phase fluctuations of oscillators can often be modelled as a simple nonrandom function of time, plus a random component. Therefore, it is useful to know the effect of certain deterministic phase functions on Allan variance.

(i) Constant phase and frequency offsets:

$$x(t) = a_0 + a_1 t$$

Here, $\Delta_\tau^2 x(t) = 0$, so $\sigma_y^2(\tau) = 0$ for all τ . Such constant offsets have no effect on the Allan variance, which is non-zero only if the frequency difference of the two oscillators is changing with time.

(ii) Linear frequency drift:

$$x(t) = \frac{1}{2} c t^2 \quad (\text{so } y(t) = ct)$$

In this case $\Delta_\tau^2 x(t) = c\tau^2$, and

$$\sigma_y(\tau) = \frac{1}{\sqrt{2}\tau} |c| \tau^2 = \frac{|c|\tau}{\sqrt{2}} \quad (6)$$

A linear frequency drift causes the Allan deviation to be proportional to τ . Section IV shows one way of estimating the drift and removing it from the measurements.

- (iii) *Higher powers of t .* If $n \geq 3$ and $x(t) = ct^n$, then

$$\Delta_t^2 x(t) = cn(n-1)\tau^2 t^{n-2} + \text{terms of lower degree}$$

Since this grows like a positive power of t , the Allan variance does not exist. (The limit in (3) is $+\infty$.)

- (iv) *A single frequency spike.* This is the same as a step in phase or time. Suppose that $x(t)$ jumps by an amount X_0 during a measurement of duration T . Given $\tau = T/m$ there is an index k such that $\Delta_\tau x(k\tau) = X_0$ and the other $\Delta_\tau x(j\tau)$ are zero. Then

$$\Delta_\tau^2 x(k\tau) = X_0, \quad \Delta_\tau^2 x((k+1)\tau) = -X_0$$

and the rest are zero. It follows from (4) that

$$S_y^2(\tau, m) = \frac{X_0^2}{\tau(T-\tau)}$$

Thus, for $\tau \ll T$, the estimated Allan variance is approximately proportional to $1/\tau$. A single frequency spike mimics white frequency noise (see below). This is appropriate, since a Poisson train of spikes is a form of white noise.

- (v) *Periodic disturbances.* These can be caused by daily temperature variations, periodic weather fluctuations, problems with the measurement system, or even problems with the oscillators themselves. Their effect on Allan variance measurements can be bizarre. For example, let

$$x(t) = X_0 \cos(2\pi\nu t + \theta)$$

Some computation yields

$$\sigma_y(\tau) =$$

$$\frac{2}{\tau} |X_0| \sin^2(\pi\nu\tau) \cdot \begin{cases} \sqrt{2} |\cos \theta|, & \text{if } 2\nu\tau = \text{integer} \\ 1, & \text{otherwise} \end{cases}$$

This is the theoretical Allan deviation (passing to the limit in (4)). For $\theta = 0$, this function looks like Fig. 3. If your Allan deviation plot has a lot of wiggles or

looks like a staircase, inspect your raw x or \bar{y} data for periodic contamination.

C. The Classical Random Clock Noise Model

This is a combination of five “power-law” random processes, each of which contributes its signature to the Allan variance. The model is specified by the one-sided spectral density $S_y(f)$ of the fractional frequency $y(t)$:

$$S_y(f) = h_2 f^2 + h_1 f + h_0 + h_{-1} f^{-1} + h_{-2} f^{-2}$$

$$= \sum_{\alpha=-2}^2 h_\alpha f^\alpha$$

Table 2 gives the Allan variances of the five components of the model (Ref. 1).

For example, if one sees $\sigma_y(\tau)$ proportional to $1/\sqrt{\tau}$ over two decades or more, one usually interprets this as the effect of white frequency noise. You had better eyeball the x or \bar{y} data versus time to make sure that the $1/\sqrt{\tau}$ is not caused by a single monster phase jump (see item (iv) above). The inference of a model from $\sigma_y(\tau)$ is unreliable without some common-sense checks.

For white phase and flicker phase, the $\sigma_y^2(\tau)$ formulas are approximations that require $2\pi f_h \tau \gg 1$. Moreover, the white phase formula can be extended to the case of a stationary time deviation process $x(t)$ having an autocovariance function $R_x(\tau) = \text{Ex}(t)x(t+\tau)$. Let $\sigma_x^2 = R_x(0)$, the time variance. If there is a τ_0 such that $|R_x(\tau)| \ll R_x(0)$ for $\tau \geq \tau_0$, then

$$\sigma_y^2(\tau) \approx \frac{3\sigma_x^2}{\tau^2}, \quad \tau \geq \tau_0 \quad (7)$$

III. Zero-Crossing Counter Technique

Figure 4 shows how the frequency and timing groups in the Communications Systems Research Section measure the time deviation $x(t)$ of two oscillators. The oscillators are offset in frequency by ν_b , which, in the present setup, is at most 1 Hz. Oscillator 2 runs at the higher frequency. If oscillators 1 and 2 have phases ϕ_1 and ϕ_2 , then the beat-note signal is

$$\sin(2\pi\nu_b t - \phi(t))$$

where $\phi(t) = \phi_1(t) - \phi_2(t)$. The zero-crossing detector triggers the counter at the beat-note upcrossing times t_0, t_1, t_2, \dots . The counter, running *continuously* at frequency ν_c (1 MHz in

the present setup), records discrete approximations t'_0, t'_1, t'_2, \dots for further digital processing. The quantization errors $t'_j - t_j$ will be dealt with later.

Let $\tau_b = 1/\nu_b$, the nominal beat period. Then $t_j - t_{j-1} \approx \tau_b$. Note that τ_b is not precisely defined by the data. It could be $t_1 - t_0$, or the average of all the $t_j - t_{j-1}$ over the duration of the test. Changing τ_b also changes $\phi(t)$, but only by $\text{const} \cdot t$, so that frequency stability is not changed.

The time deviation of the *beat note* is

$$x_b(t) = -\frac{\phi(t)}{2\pi\nu_b} = -\frac{\nu_0}{\nu_b} \frac{\phi(t)}{2\pi\nu_0} = -\frac{\nu_0}{\nu_b} x(t)$$

where $x(t)$ is the time deviation of the oscillator pair. (Note the sign reversal.) Since the factor ν_0/ν_b is large, typically 10^6 to 10^8 , the fluctuations of the oscillators are magnified so that they can be measured. The Allan deviation of the oscillators is ν_b/ν_0 times the Allan deviation of the beat note.

The computation of the Allan variance of the beat note is not totally straightforward. By definition, the j^{th} upcrossing t_j occurs when the total phase of the beat note is $2\pi j$, i.e.,

$$2\pi\nu_b t_j - \phi(t_j) = 2\pi j$$

At time t_j the time deviation of the beat note is

$$x_b(t_j) = -\frac{\phi(t_j)}{2\pi\nu_b} = j\tau_b - t_j, \quad j = 0, 1, 2, \dots$$

(another sign reversal), and hence the time deviation of the oscillators is

$$x(t_j) = \frac{\nu_b t_j - j}{\nu_0}$$

Now suppose that the aim is to measure $\sigma_y(\tau)$ for $\tau = r\tau_b$, r a positive integer. This could be done if we had the sequence

$$x(t_0), x(t_0 + \tau), x(t_0 + 2\tau), \dots \quad (8)$$

What we *do* have is the sequence

$$x(t_0), x(t_r), x(t_{2r}), \dots \quad (9)$$

which is x at times that are not exactly spaced by τ . This is all right as long as the periods $t_{jr} - t_{(j-1)r}$ differ from τ by at most 1%, say. In other words, the fractional frequency stabil-

ity of the beat note must be no worse than 1%. The beat note itself must be a reasonably good clock, but not so good that its fluctuations are hidden by counter quantization. In most of our tests, this stability is 10^{-4} or better. There was one test, however, in which the beat periods $t_j - t_{j-1}$ fluctuated by a factor of two or more. In this situation, the Allan variance estimate given below is invalid. It is thought, however, that there might be a more sophisticated algorithm that would still be able to estimate Allan variance for $\tau \gg \tau_b$.

Since the second differences of j are zero, we have

$$\Delta_r^2 x(t_{jr}) = \frac{\nu_b}{\nu_0} \Delta_r^2 t_{jr} = \frac{\nu_b}{\nu_0} (t_{jr} - 2t_{(j-1)r} + t_{(j-2)r})$$

Using this in (4) in place of $\Delta_r^2 x(j\tau)$ gives the Allan variance estimate

$$S_y^2(\tau, m) = \frac{\nu_b^2}{\nu_0^2} \frac{1}{2\tau^2} \frac{1}{m-1} \sum_{j=2}^m [\Delta_r^2 t_{jr}]^2$$

where $\tau = r\tau_b$ and the duration of the test is $m\tau$. Look again at the numerical example in Section II. The $x(t)$ column of Table 1 actually gives t'_{jr} in microseconds ($r = 256, j = 0$ to $8, \tau_b = 1$ s) with the gross linear part $j r \tau_b$ already subtracted off. Since $\nu_b = 1$ Hz, $\nu_0 = 10^8$ Hz, the scale factor is $10^{-6} \nu_b/\nu_0 = 10^{-14}$.

Quantization Error. Let the counter frequency be ν_c . (Ours is 1 MHz.) My model for the recorded time t'_j of the j^{th} upcrossing t_j is

$$t'_j = t_j + q_j$$

where t'_j is a multiple of $1/\nu_c$, and q_j , the quantization error, satisfies $0 \leq q_j < 1/\nu_c$. As an approximation, assume that the processes t_j and q_j are uncorrelated, and that q_j is a process of independent, uniform random variables. This approximation also yields Sheppard's second-moment correction in statistics. Then the q_j alone contribute an Allan variance

$$\sigma_q^2(\tau) = \left(\frac{\nu_b}{\nu_0} \frac{1}{2\nu_c \tau} \right)^2$$

(This follows from (7).) For example, if $\nu_0 = 5$ MHz, $\nu_b = 1$ Hz, $\nu_c = 1$ MHz, then $\sigma_q(\tau) = 10^{-13}/\tau$. One should *subtract* $\sigma_q^2(\tau)$ from the measured $\sigma_y^2(\tau)$ (from t'_j) to get the true $\sigma_y^2(\tau)$ (from t_j). In other words, $\sigma_q^2(\tau)$ contributes to the measurement-system noise floor.

IV. Drift Removal

During a long test of a pair of hydrogen masers, it is often apparent that the frequency deviation $\nu(t)$ appears to be dominated by a linear component, which also causes $\sigma_{\nu}(\tau)$ to increase like τ for large τ , thus masking the effect of the random fluctuations. See Fig. 5 and the upper curve of Fig. 6. The model is

$$x(t) = a_0 + a_1 t + \frac{1}{2} c t^2 + x_0(t)$$

or

$$\nu(t) = a_1 + c t + \frac{dx_0}{dt}$$

where c is the frequency drift rate ($\Delta f/f$ per second) and $x_0(t)$ is a mean-zero process. The constants a_0, a_1 are irrelevant for a frequency stability study. Let $\sigma_{\nu}^2(\tau), \sigma_{x_0}^2(\tau)$ be the Allan variances of $x(t), x_0(t)$. We would like to estimate c and to use the estimate \hat{c} to extract an estimate of $\sigma_{x_0}^2(\tau)$, the Allan variance of the underlying fluctuations.

A. Estimating \hat{c} From $x(t)$, $0 \leq t \leq T$

Assume a sample time of τ_0 . The \hat{c} estimators given below are all unbiased.

Method 0. Least-squares quadratic fit to $x(t)$:

$$\hat{x}(t) = \hat{a}_0 + \hat{a}_1 t + \frac{1}{2} \hat{c}_0 t^2$$

where $\hat{a}_0, \hat{a}_1, \hat{c}_0$ are chosen to minimize $\sum [\hat{x}(t) - x(t)]^2$. This method is optimal (minimal variance) in the presence of white phase noise.

Method 1. Least-squares linear fit to $\bar{\nu}(t, \tau_0)$:

$$\hat{\nu}(t) = \hat{a}_1 + \hat{c}_1 t$$

where \hat{a}_1, \hat{c}_1 are chosen to minimize $\sum [\hat{\nu}(t) - \bar{\nu}(t, \tau_0)]^2$. This method is optimal in the presence of white frequency noise.

Method 2. Mixed second difference of phase. This is the one that has been used here for tests of frequency standards (Refs. 4 and 7). No one else is thought to have used it. So far, no complaints have been received. Let

$$\hat{c}_2 = \frac{\bar{\nu}(T, \tau_c) - \bar{\nu}(\tau_c, \tau_c)}{T - \tau_c}$$

In other words, drift rate equals average frequency at the end of the run, minus average frequency at the beginning of the run, divided by the time span between the midpoints of the averaging intervals $[0, \tau_c]$ and $[T - \tau_c, T]$. Take

$$\tau_c = \frac{T}{6.29}$$

a value that minimizes the variance of \hat{c}_2 in the presence of flicker-frequency noise. This value is not critical. In terms of $x(t)$,

$$\hat{c}_2 = \frac{x(T) - x(T - \tau_c) - x(\tau_c) + x(0)}{\tau_c(T - \tau_c)} \quad (10)$$

so that only four values of $x(t)$ are needed.

This estimator, although suboptimal for all the classical noises, performs well in the presence of white, flicker, and random walk frequency noises. If $x_0(t)$ is random walk frequency, then the standard deviation of \hat{c}_2 is

$$\sigma(\hat{c}_2) = 4.6 \frac{\sigma_{x_0}(\tau_c)}{T} \quad (11)$$

and this formula is pessimistic for white and flicker frequency. Note that you have to remove drift in order to estimate $\sigma_{x_0}(\tau_c)$; see (13) below.

B. Removing Drift

Suppose that some unbiased estimate \hat{c} is obtained, perhaps one of the foregoing, perhaps some other. An estimate of the residual $x_0(t)$ is

$$\hat{x}_0(t) = x(t) - \frac{1}{2} \hat{c} t(t - T)$$

See Fig. 7 for the effect of this ($\hat{c} = \hat{c}_2$) on the data of Fig. 5. Then

$$\Delta_{\tau}^2 \hat{x}_0(t) = \Delta_{\tau}^2 x(t) - \hat{c} t^2$$

is used in the estimate

$$S_{y_0}^2(\tau, m) = \frac{1}{2\tau^2} \frac{1}{m-1} \sum_{j=2}^m \left[\Delta_{\tau}^2 \hat{x}_0(j\tau) \right]^2$$

for $\sigma_{y_0}^2(\tau)$. Expanding the square gives

$$S_{y_0}^2(\tau, m) = S_y^2(\tau, m) - \frac{\tau^2}{2} c_{\tau}^2 + \frac{\tau^2}{2} (c_{\tau} - \hat{c})^2 \quad (12)$$

where

$$c_\tau = \frac{x(T) - x(T-\tau) - x(\tau) + x(0)}{\tau(T-\tau)}$$

and $T = m\tau$. If $\hat{c} = \hat{c}_2$, then

$$S_{y0}^2(\tau_c, 6) = S_y^2(\tau_c, 6) - \frac{\tau_c^2}{2} \hat{c}_2^2 \quad (13)$$

is used for estimating $\sigma_{y0}^2(\tau_c)$, needed in (11). The lower curve in Fig. 6 is $S_{y0}(\tau, m)$, the upper curve is $S_y(\tau, m)$, and the straight line is (6) with $c = \hat{c}_2$.

CAUTION

If \hat{c} is based on the test data $x(t)$ ($0 \leq t \leq T$), $T = m\tau$, m is small (≤ 5 , say), and random walk frequency noise dominates $x_0(t)$, then $S_{y0}^2(\tau, m)$ severely underestimates $\sigma_{y0}^2(\tau)$ on the average, and has a larger variance than $S_y^2(\tau, m)$. If $\hat{c} = \hat{c}_2$, the bias of S_{y0}^2 is given by Ref. 7.

m	10	7	5	4	3	2
Bias	-16%	-24%	-34%	-43%	-59%	-89%

The error bars in Fig. 6 are the naive ones, and do not reflect these biases. The lesson here is to use a long-term estimate of c if possible. For example, cavity retuning of a hydrogen maser gives an independent measurement of frequency drift.

V. Three-Cornered Hat

The previous material has been about the relative stability of a pair of oscillators. If one has a triplet of oscillators, one may be able to estimate the stability of each one individually by using the three pair-comparisons. The setup, as shown in Fig. 8, has oscillators A, B, C, and pair-channels 1, 2, 3. We measure the pair Allan variances $\sigma_1^2(\tau)$, $\sigma_2^2(\tau)$, $\sigma_3^2(\tau)$ and would like to compute $\sigma_A^2(\tau)$, $\sigma_B^2(\tau)$, $\sigma_C^2(\tau)$ for the individual oscillators.

Assumptions

- (1) The phase fluctuations of the three oscillators are independent.

- (2) Relative linear frequency drift is negligible or removed. In other words, the second differences of x_A , x_B , x_C have mean zero.

If the assumptions are satisfied, then, for a fixed τ , the processes $\Delta_\tau^2 x_A$, $\Delta_\tau^2 x_B$, $\Delta_\tau^2 x_C$ are orthogonal. Therefore,

$$\begin{aligned} \sigma_1^2 &= \sigma_B^2 + \sigma_C^2 \\ \sigma_2^2 &= \sigma_C^2 + \sigma_A^2 \\ \sigma_3^2 &= \sigma_A^2 + \sigma_B^2 \end{aligned} \quad (14)$$

and so

$$\begin{aligned} \sigma_A^2 &= \frac{1}{2} (\sigma_2^2 + \sigma_3^2 - \sigma_1^2) \\ \sigma_B^2 &= \frac{1}{2} (\sigma_3^2 + \sigma_1^2 - \sigma_2^2) \\ \sigma_C^2 &= \frac{1}{2} (\sigma_1^2 + \sigma_2^2 - \sigma_3^2) \end{aligned} \quad (15)$$

Figure 8 shows the geometric interpretation.

Difficulties

- (1) Since frequency standards are sensitive to changes in the environment (temperature, pressure, humidity, magnetic field), it seems that frequency standards in the same room must violate assumption 1 for long-term fluctuations, say, for 1000 s or more.
- (2) The pair-channel $\sigma^2(\tau)$ values are not available, just estimates such as $S^2(\tau, m)$. When estimates are substituted for the σ^2 values in the right side of (15), it often happens that some oscillator σ^2 comes out negative. One interpretation of this is that a confidence interval for that σ goes all the way down to zero. For example, if $S_1^2 > S_2^2 + S_3^2$, then the σ_A^2 estimate is negative. (If a triangle has sides S_1 , S_2 , S_3 , then the angle between S_2 and S_3 is obtuse.) If the stronger inequality $S_1 > S_2 + S_3$ holds, then something is wrong with the measurements. (There is no triangle with sides S_1 , S_2 , S_3 .)

References

1. Rutman, J., "Characterization of phase and frequency instabilities in precision frequency sources: fifteen years of progress," *Proc. IEEE*, Vol. 66, pp. 1048-1075, 1978.
2. Lesage, P., and Audoin, C., "Estimation of the two-sample variance with limited number of data," *Proc. 31st Ann. Symp. on Freq. Control*, Document No. ADA088221, National Technical Information Service, Springfield, VA, pp. 311-318, 1977.
3. Yoshimura, K., "Characterization of frequency stability: uncertainty due to the autocorrelation of the frequency fluctuations," *IEEE Trans. Instrum. Meas.*, Vol. IM-27, pp. 1-7, 1978.
4. Greenhall, C., "Removal of drift from frequency stability measurements," *TDA Progress Report 42-65*, pp. 127-132, Jet Propulsion Laboratory, Pasadena, California, 1981.
5. Greenhall, C., "A structure function representation theorem with applications to frequency stability estimation," *IEEE Trans. Instrum. Meas.*, Vol. IM-32, pp. 364-370, 1983.
6. Percival, D., *The Statistics of Long Memory Processes*, Ph.D. Thesis, U. of Washington, 1983.
7. Greenhall, C., "The fundamental structure function of oscillator noise models," *Proc. 14th Ann. PTTI Meeting* (NASA Conference Publication 2265), pp. 281-294, 1982.
8. Howe, D., et al., "Properties of signal sources and measurement methods," *Proc. 35th Ann. Freq. Control. Symp.*, Electronic Industries Association, Washington, DC, pp. A1-A47, 1981.

Table 1. Data from a stability test of two hydrogen masers

t	$x(t)$	$\Delta_1 x(t)$	$\Delta_1^2 x(t)$	$\Delta_2 x(t)$	$\Delta_2^2 x(t)$	$\Delta_3 x(t)$	$\Delta_3^2 x(t)$
0	0						
1	658	658					
2	1229	571	-87	1229			
3	1701	472	-99			1701	
4	2333	632	160	1104	-125	1675	
5	2991	658	26			1762	
6	3493	502	-156	1160	56	1792	91
7	4095	602	100			1762	87
8	4690	595	-7	1197	37	1699	-63

Table 2. Spectral density and Allan variance of the components of the classical clock noise model

α	Name	$S_y(f)$	$\sigma_y^2(\tau)$
2	White phase	$h_2 f^2, f < f_h$	$h_2 f_h \frac{3}{4\pi^2 \tau^2}$
1	Flicker phase	$h_1 f, f < f_h$	$h_1 \frac{3}{4\pi^2 \tau^2} \ln(8.88 f_h \tau)$
0	White frequency	$h_0, f_h = +\infty$	$h_0 \frac{1}{2\tau}$
-1	Flicker frequency	$h_{-1} f^{-1}, f_h = +\infty$	$h_{-1} \ln 4$
-2	Random walk frequency	$h_{-2} f^{-2}, f_h = +\infty$	$h_{-2} \frac{2\pi^2 \tau}{3}$




Fig. 1. x and \bar{y} from a test of hydrogen masers




Fig. 2. Estimated Allan deviation of a pair of hydrogen masers




Fig. 3. Allan deviation for $x(t) = \cos(2\pi\nu t)$

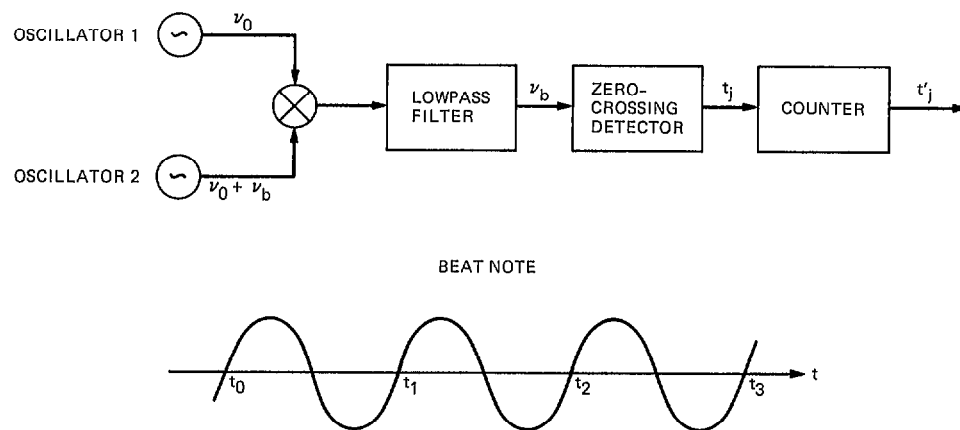


Fig. 4. Frequency stability test setup




Fig. 5. x and \bar{y} before drift removal




Fig. 6. Estimated Allan deviation before and after drift removal




Fig. 7. x and \bar{y} after drift removal




Fig. 8. Three-cornered hat